arXiv:1605.09520v1 [cs.DS] 31 May 2016

Simpler Self-reduction Algorithm for Matroid Path-width

Petr Hliněný*

Faculty of Informatics, Masaryk University Brno, Czech Republic
hlineny@fi.muni.cz

Abstract

Path-width of matroids naturally generalizes better known path-width of graphs, and is NP-hard by a reduction from the graph case. While the term *matroid path-width* was formally introduced by Geelen–Gerards–Whittle [JCTB 2006] in pure matroid theory, it was soon recognized by Kashyap [SIDMA 2008] that it is the same concept as long-studied so called trellis complexity in coding theory, later named *trellis-width*, and hence it is an interesting notion also from the algorithmic perspective. It follows from a result of Hliněný [JCTB 2006], that the problem to test whether a given matroid over a finite field has path-width at most t is fixed-parameter tractable in t, but this result does not give any clue about a corresponding path-decomposition.

The first constructive, though rather complicated, FPT algorithm for path-width of matroids over a finite field has been given just recently by Jeong–Kim–Oum [SODA 2016]. Here we give a much simpler *self-reduction* constructive FPT algorithm for the same problem. Precisely, we design an efficient routine that constructs an optimal path-decomposition of a matroid (even an abstract one, given by a rank oracle) using a "black-box" subroutine for testing whether the path-width of a matroid is at most t. In connection with the aforementioned decision algorithm for path-width of matroids over a finite field we then get the desired constructive FPT algorithm.

1998 ACM Subject Classification (2012) Theory of computation → Fixed parameter tractability; Mathematics of computing → Matroids, Coding theory

Keywords and phrases matroid; path-width; trellis-width; fixed-parameter tractability; well-quasi-ordering

1 Introduction

An ordinary path-decomposition of a graph G, see [15], is a sequence of sets $\big(X_i \subseteq V(G) : i = 1,\dots,p\big)$, such that; (i) $\bigcup_{i=1}^{p} X_i = V(G)$ and for every $1 \le i < j < k \le p$, it is $X_j \subseteq X_i \cap X_k$, and (ii) for every $e = uv \in E(G)$ there is $1 \le i \le p$ such that $u, v \in X_i$. The width of the decomposition equals $\max_{1\le i\le p} |X_i| - 1$, and the *path-width* of G is the minimum width over all path-decompositions of G. This notion, together with related tree-width, have received great attention in the Graph Minors project of Robertson and Seymour.

There is another, more recent view of path-width; the *matroid path-width* defined first by Geelen, Gerards and Whittle [4] in matroid research. We refer to Section 2 for the definition. While the two variants of path-width are indeed tightly related, there is no simple explicit formula between the ordinary path-width and the matroid path-width of the same graph. Matroid path-width of graphs has been recently studied in some papers, e.g. [12]. Our interest in matroid path-width, however, lies beyond the graph case.

A similar notion to path-width has been considered for quite some time also in the area of coding theory, under various names such as the "trellis complexity" of a code, e.g. [9, 16].

* This research was supported by the Czech Science Foundation project No. 14-03501S.

© Petr Hliněný;
licensed under Creative Commons License CC-BY
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In 2008, Kashyap [11] observed that this is the same parameter as the aforementioned path-width [4] of a vector matroid represented by the generator matrix of a linear code. He introduced for it the new name *trellis-width* of a linear code, and proved that computing trellis-width is NP-hard by a reduction from graph path-width. Kashyap also asked, as one of the main open problems in [11], how difficult is to decide whether the trellis-width of a linear code is at most t, and to construct the corresponding optimal decomposition in the YES case, where $t \in \mathbb{N}$ is a fixed parameter.

Concerning the first half of Kashyap's question, the decision problem is in FPT (*fixed-parameter tractable*) which follows already from earlier author's papers [5, 7]. Recall that a parameterized problem is in FPT if it admits an algorithm with runtime of order $\mathcal{O}(f(t)\cdot n^c)$ where t is the parameter, n the input size and c a constant. Let us now briefly sketch the two key ideas on which an FPT algorithm for deciding 'trellis-width $\leq t$' can be based (see Section 4 for full details):

- The branch-width of the underlying vector matroid of a code is upper-bounded in terms of t, the assumed trellis-width bound. Hence there are only finitely many "minimal obstructions" for the property 'trellis-width $\leq t$' for each $t \in \mathbb{N}$, which follows from [3]. (A similar observation occurs also in Kashyap [11].)
- For bounded branch-width of a vector matroid, over any finite field, we can construct an approximate branch-decomposition of it in FPT, see [6]. Then we can, again in FPT, check presence of each one of these finitely many obstructions, see [5, 7].

A careful reader may immediately notice a problem of the suggested scheme—in what way can we get a corresponding trellis- or path-decomposition from it? The sad truth is that in *no way*. To get a corresponding decomposition, a new approach is needed.

The problem is analogous to that of constructing an optimal matroid branch-decomposition, for which the aforementioned paper [6] provided an approximate construction and an exact value in FPT. Building upon that, Oum and the author [8] later designed a self-reduction routine which constructs an optimal branch-decomposition of a matroid over a finite field, by calling the decision subroutine for exact branch-width. It appears very natural to try the same approach also for path-decompositions, but this, unfortunately, does not easily work. Instead, just recently, Jeong, Kim and Oum [10] designed an involved standalone algorithm for the construction of an optimal path-decomposition of a matroid over a finite field, which runs in FPT time for the parameter path-width. In their algorithm, they refer back to the ideas and techniques of Bodlaender and Kloks [1] for graphs.

We, on the other hand, provide a new self-reduction approach to the matroid path-decomposition problem, partially inspired by [8]. Its main advantage over [10] is a much simpler design and a relatively short proof. Precisely, we give the following:

- A non-uniform FPT algorithm that, for fixed parameters t and $|\mathbb{F}|$, inputs an n-element matroid M represented by a matrix over a finite field $\mathbb{F}$, and in $\mathcal{O}(n^3)$ time constructs a path-decomposition of M of width $\leq t$ or concludes that the path-width of M is $> t$. *(Section 3 – Theorem 3.2, and Theorem 4.3 for improved runtime)*
- An FPT algorithm that, for a fixed parameter t, a given *oracle function* $\mathcal{P}$ testing if the path-width of a matroid is $\leq t$, and an input n-element abstract matroid M, constructs a path-decomposition of M of width $\leq t$ or concludes that the path-width of M is $> t$. *(Section 4 – Theorem 4.4)*
- A non-uniform FPT algorithm that, for a fixed parameter t, inputs an n-vertex graph G, and in $\mathcal{O}(n^3)$ time constructs a linear rank-decomposition of G of width $\leq t$ or concludes that the linear rank-width of G is $> t$. *(Section 5)*

2 Preliminaries

We refer to the standard textbook Oxley [14] for general matroid material and terminology.

Matroids; rank and connectivity. A *matroid* is a pair $M = (E, \mathcal{B})$ where $E = E(M)$ is the ground set of M (elements of M), and $\mathcal{B} \subseteq 2^E$ is a nonempty collection of *bases* of M, no two of which are in an inclusion. Moreover, matroid bases satisfy the "exchange axiom": if $B_1, B_2 \in \mathcal{B}$ and $x \in B_1 \setminus B_2$, then there is $y \in B_2 \setminus B_1$ such that $(B_1 \setminus \{x\}) \cup \{y\} \in \mathcal{B}$. We consider only finite matroids.

All matroid bases have the same cardinality called the *rank* $r(M)$ of the matroid. Subsets of bases are called *independent*, and sets that are not independent are *dependent*. Minimal dependent sets are called *circuits*. The *rank function* $r_M(X)$ of M maps subsets of $E(M)$ to non-negative integers; $r_M(X)$ equals the maximum cardinality of an independent subset of a set $X \subseteq E(M)$. The rank function is *submodular*, meaning that $r_M(X) + r_M(Y) \geq r_M(X \cup Y) + r_M(X \cap Y)$ for any $X, Y \subseteq E(M)$, and it fully defines a matroid on its ground set. A matroid M is *uniform* if all subsets of $E(M)$ of size equal $r(M)$ are bases, and it is also denoted by $U_{r,n}$ where $r = r(M)$ and $n = |E(M)|$.

For $X \subseteq E$, *deletion of* X results in the matroid $M \setminus X$ which is defined by the restriction of the rank function r_M to $E \setminus X$. On the other hand, *contraction of* X results in the matroid M/X which is defined by the rank function $r_{M/X}(Y) := r_M(X \cup Y) - r_M(X)$ for all $Y \subseteq E \setminus X$. Matroids of the form $M/X \setminus Y$ are called *minors* of M.

In a folklore example, the *cycle matroid of a graph* G is the matroid $M(G)$ on the ground set $E(G)$ such that the independent sets of $M(G)$ are acyclic edge subsets of G, circuits are ordinary cycles, the rank function equals number of vertices minus number of components, and deletion and contraction correspond to ordinary edge deletion and contraction in G.

We, moreover, define the (symmetric and submodular) *connectivity function* of M by

$$\lambda_M(X) := r_M(X) + r_M(E \setminus X) - r(M)$$

for all subsets $X \subseteq E$. Any bipartition (X, Y) of E (where $Y = E \setminus X$) is called a *separation in* M of connectivity value $\lambda_M(X) = \lambda_M(Y)$, or shortly a k-separation if $\lambda_M(X) = k - 1$. Informally, λ_M measures how much the two sides of a separation "share together" in terms of rank. We also define

$$\mu_M(X, A) := r_M(X \cup A) + r_M\big((E \setminus X) \cup A\big) - r(M),$$

which, very informally saying, "adds" A (the rank of) to the connectivity value of $(X, E \setminus X)$.

The *closure of a set* $X \subseteq E$ *in* M, denoted by $cl_M(X)$, is defined by

$$cl_M(X) := \big\{e \in E : r_M(X \cup \{e\}) = r_M(X)\big\}.$$

The closure of X, hence, includes all elements dependent on (or spanned by) X. Sets X such that $X = cl_M(X)$ are closed, or *flats*. For example, $\mu_M(X, A) = \lambda_M(X)$ if and only if $A \subseteq cl_M(X) \cap cl_M(E \setminus X)$ or, in other words, if A is spanned by both X and $E \setminus X$. If $e \in cl_M(X) \cap cl_M(E \setminus X)$, then we say that e is *in the guts of the separation* $(X, E \setminus X)$.

Matroid path-width. Let M be an n-element matroid. Any permutation $(e_1, e_2, \ldots, e_n)$ of the elements $E(M)$ is called a *path-decomposition* of M. The *width of* $(e_1, e_2, \ldots, e_n)$ is defined

$$w_M(e_1, e_2, \ldots, e_n) := \max_{i=1,\ldots,n} \lambda_M(\{e_1, \ldots, e_i\}),$$

and the *path-width* $\mathrm{pw}(M)$ *of* M is the least width over all path-decompositions of M, i.e.

$$\mathrm{pw}(M) := \min_{\text{permut. } \pi \in S_n} w_M\left(e_{\pi(1)}, e_{\pi(2)}, \ldots, e_{\pi(n)}\right).$$

We say, for any $1 \le i < n$, that the separation $\big(\{e_1, \ldots, e_i\}, \{e_{i+1}, \ldots, e_n\}\big)$ is *displayed* by the path-decomposition $(e_1, e_2, \ldots, e_n)$, and we refer to $\big(\{e_1, \ldots, e_i\}, \{e_{i+1}, \ldots, e_n\}\big)$ as to the *separation at position* i.

The notion of matroid path-width is related to better known branch-width. A tree T is cubic if its vertex degrees are 3 or 1. A *branch-decomposition* of a matroid M is a pair (T, τ) where T is a cubic tree and $\tau : E(M) \to \ell(T)$ is a bijection of the elements of M to the leaves of T. Every edge $e \in E(T)$ partitions the leaves of T into two sets L_1, L_2, and we say the separation $(\tau^{-1}(L_1), \tau^{-1}(L_2))$ is *displayed by* (T, τ). We define the width of e as $\lambda_M(\tau^{-1}(L_1)) + 1$ and the width of (T, τ) as the maximum of widths over all edges of T. The *branch-width* $\mathrm{bw}(M)$ of M is the minimum width over all branch-decompositions of M.

A cubic tree is a caterpillar if it is obtained by connecting leaves to a path. *Linear branch-width* of a matroid M is defined as ordinary branch-width with a restriction that the cubic tree T must be a caterpillar. One can easily observe that this notion coincides with that of matroid path-width (except the artificial '+1' term above); the path-width of M is always one less than its linear branch-width. Consequently, we have:

▶ **Lemma 2.1.** *For any matroid* M*, it is* $\mathrm{bw}(M) \le \mathrm{pw}(M) + 1$. ◀

Assorted claims. The following routine technical claims about matroids will be used in the proof of our algorithm.

▶ **Lemma 2.2.** *Let* M *be a matroid and* $C_1, C_2 \subseteq E(M)$ *be two circuits of* M *such that* $|C_1 \cap C_2| = 1$ *and* $r_M(C_1) + r_M(C_2) = r_M(C_1 \cup C_2) + 1$*. Then* $C_1 \Delta C_2$ *(the symmetric difference) is also a circuit of* M*.*

Proof. Let $C_1 \cap C_2 = \{f\}$. By the folklore circuit exchange axiom there exists a circuit of M contained in the set $C_3 := (C_1 \cup C_2) \setminus \{f\} = C_1 \Delta C_2$. Consider any $e \in C_3$ where, up to symmetry, $e \in C_2 \setminus C_1$. It is $|C_3 \setminus \{e\}| = |C_1| - 1 + |C_2| - 1 - 1 = r_M(C_1) + r_M(C_2) - 1 = r_M(C_1 \cup C_2)$. At the same time, since C_1, C_2 are circuits and $f \in C_1 \cap C_2$, it is $r_M(C_1 \cup C_2) = r_M\big(C_1 \cup (C_2 \setminus \{e\})\big) = r_M\big((C_1 \setminus \{f\}) \cup (C_2 \setminus \{e, f\})\big) = r_M(C_3 \setminus \{e\})$ and so $C_3 \setminus \{e\}$ is independent. Therefore, C_3 itself is the circuit. ◀

▶ **Lemma 2.3.** *Let* M *be a matroid and* $X \subseteq E = E(M)$*. If* $e, f \in E$ *such that* $\mu_M(X, \{e\}) = \mu_M(X, \{f\}) = \mu_M(X, \{e, f\}) = \lambda_M(X) + 1$*, then either* $e, f \in X$ *or* $e, f \notin X$*.*

Proof. Let $Y = E \setminus X$. Assume the contrary, i.e. up to symmetry, $e \in X$ and $f \in Y$. From $r_M(X) + r_M(Y) - r(M) + 1 = \lambda_M(X) + 1 = \mu_M(X, \{f\}) = r_M(X \cup \{f\}) + r_M(Y) - r(M)$ we immediately get $r_M(X \cup \{f\}) = r_M(X) + 1$ and, by symmetry, $r_M(Y \cup \{e\}) = r_M(Y) + 1$. This leads to

$$\begin{aligned} \lambda_M(X) + 1 &= \mu_M(X, \{e, f\}) = r_M(X \cup \{f\}) + r_M(Y \cup \{e\}) - r(M) \\ &= r_M(X) + r_M(Y) + 2 - r(M) = \lambda_M(X) + 2, \end{aligned}$$

a contradiction. ◀

▶ **Lemma 2.4.** *Let* M *be a matroid and* N *a minor of* M*. Then* $\mathrm{pw}(N) \le \mathrm{pw}(M)$*.*

▶ **Lemma 2.5.** *Let M be an n-element matroid and $(e_1, \ldots, e_n)$ be its path-decomposition of width $t = w_M(e_1, \ldots, e_n)$. For an index i let $X = \{e_1, \ldots, e_i\}$ and $Y = \{e_{i+1}, \ldots, e_n\} = E(M) \setminus X$ such that $\lambda_M(X) = t$. Assume that there exists a circuit $C \subseteq E(M)$ such that no element of C is in the guts of (X, Y) and $\mu_M(X, C) = r_M(C)$. Then $X \cap C \neq \emptyset \neq Y \cap C$.*

(The condition $\mu_M(X, C) = r_M(C)$ can be seen as that "C spans the guts of (X, Y).")

Due to restricted space in this short paper we leave the proofs out (see the Appendix).

Matroid representation and extensions. A standard example of a matroid is given by a set of vectors (forming the columns of a matrix $\boldsymbol{A}$) with usual linear independence. The matrix $\boldsymbol{A}$ is then called a *(vector) representation* of the matroid. We will consider only representations $\boldsymbol{A}$ over finite fields. Since non-zero scaling of vectors does not change linear dependencies, vector representations can also be seen as point configurations in the projective space over $\mathbb{F}$ (note that parallel vectors are represented by identical points).

We now briefly illustrate the "geometric" meaning of matroid terms.

- The matroid closure of a set X corresponds to the affine closure or *span* $\langle X \rangle$ of the points representing X. The rank of X is the dimension or rank of the span of X.
- For a separation (X, Y) of M, the *guts* of (X, Y) corresponds to the intersection of the spans of X and Y, and $\lambda_M(X)$ is the rank of this guts. The value of $\mu_M(X, A)$ equals the rank of the space spanned by the guts $\langle X \rangle \cap \langle Y \rangle$ and A together.
- All the previous entities can be straightforwardly computed by means of standard linear algebra over the matrix $\boldsymbol{A}$.

There is one particular operation we need to discuss in close details. For a matroid M we say that a matroid M_1 is a *free extension of M by element e* if $e \in E(M_1)$ and $M = M_1 \setminus e$, $r(M_1) = r(M)$, and for every $X \subseteq E(M)$ it is $r_{M_1}(X \cup \{e\}) = r_M(X) + 1$ unless $r_M(X) = r(M)$. This is equivalent to claiming that every circuit of M_1 containing e has full rank $r(M_1)$. Informally saying, e is added to M without any unforced dependency (geometrically, in a general position). We will use the following folklore fact:

▶ **Lemma 2.6.** *Let M be a matroid of rank r represented by a matrix $\boldsymbol{A}$ over a finite field $\mathbb{F}$. Let α be a root of an irreducible polynomial of degree r in $\mathbb{F}$, and denote by $\boldsymbol{b} = (1, \alpha, \ldots, \alpha^{r-1})^T$. Let $\mathbb{F}(\alpha)$ be the extension field of $\mathbb{F}$ obtained by adjoining α to $\mathbb{F}$. Then the matrix $[\boldsymbol{A}|\boldsymbol{b}]$ over $\mathbb{F}(\alpha)$ represents a free extension of M by an element b.*

Proof. Assume the contrary, that $\boldsymbol{b}$ is a linear combination over $\mathbb{F}(\alpha)$ of the columns of a column-submatrix $\boldsymbol{A'} \subseteq \boldsymbol{A}$ of rank less than r. Since $\boldsymbol{A'}$ has r rows denoted by $\boldsymbol{a}'_1, \boldsymbol{a}'_2, \ldots, \boldsymbol{a}'_r$, they are linearly dependent as vectors, and so for some $\lambda_1, \ldots, \lambda_r \in \mathbb{F}$ (not all 0) it holds $\lambda_1 \boldsymbol{a}'_1 + \lambda_2 \boldsymbol{a}'_2 + \ldots, \lambda_r \boldsymbol{a}'_r = \boldsymbol{0}$. However, since $\boldsymbol{b}$ is a linear combination of the columns of $\boldsymbol{A'}$, it is also $\lambda_1 + \lambda_2 \alpha^1 + \ldots, \lambda_r \alpha^{r-1} = 0$. This contradicts the assumption that α is a root of an irreducible polynomial of degree r over $\mathbb{F}$. ◀

3 Self-reduction Algorithm

The purpose of this core section of the paper is to design and prove a self-reduction routine that, for a fixed parameter t, constructs an optimal path-decomposition of a given represented matroid of path-width t, using an oracle which can test whether the path-width of any given represented matroid is at most t.

Prologue. For better understanding, we start this section with a brief overview of the key part of the algorithm for constructing a matroid branch-decomposition of optimal width t from [8], which has been an inspiration for our algorithm. The key decision step in its recursive self-reduction routine is as follows:

- Assume $X \subseteq E = E(M)$ is such that $\lambda_M(X) \leq t$ and that $M[X]$ (the restriction of M to X) has branch-width $\leq t$. The task is to decide whether M has a branch-decomposition of width t such that the separation $(X, E \setminus X)$ is displayed by it (i.e., informally, that "X sticks together" as one branch of the decomposition).

In the overview, we treat a given matroid M represented over a finite field $\mathbb{F}$ as a point configuration in a projective geometry over $\mathbb{F}$. Recall from Section 2 that a point (matroid element) e is in the guts of a separation (X, Y) if e belongs to both the spans (closures in matroid sense) of X and of Y. So, in our geometric setting, we may rigorously refer to the subspace $\langle X\rangle \cap \langle Y\rangle$, which is the intersection of the (geometric) spans of X, Y, as to the *guts of* (X, Y). Therefore, the above question may be reformulated as (*) whether there exists a branch-decomposition of width $\leq t$ of any matroid $M' \supseteq M \setminus X$, which displays the guts of $(X, E \setminus X)$; if so, then a branch-decomposition of $M[X]$ can be "glued" to this place.

The solution to (*) given in [8] replaces X in M with a copy U of the uniform matroid $U_{k,3k-2}$, where $k = \lambda_M(X)$, such that $\langle U\rangle$ makes the guts of $(X, E \setminus X)$. Let M' denote the resulting matroid. One can argue that the answer to (*) is logically equivalent to whether the branch-width of M' is at most t. In one direction, it is trivial to extend any branch-decomposition displaying $\langle U\rangle$ with a new branch decomposing U. In the other direction, by the definition of a uniform matroid, every k-tuple of points of U is a basis of $\langle U\rangle$ and every partition of U into three parts has one of size at least k. Consequently, in every branch-decomposition of width $\leq t$ of M' there is an edge displaying a separation whose guts contains a basis of U, and hence displaying desired $(X, E \setminus X)$ back in M.

We remark that, in general, the matroid U may not be representable over the field $\mathbb{F}$. In such a case we use a suitable finite extension field of $\mathbb{F}$ instead (cf. Lemma 2.6).

Our aim is to design an analogous self-reduction routine for matroid path-decomposition, that is, one based on a decision step asking if a matroid has an optimal path-decomposition displaying the specified prefix. Our new solution follows.

Algorithm outline. We now give a high-level description of our new path-decomposition algorithm. Let M be the input matroid and $E = E(M)$, $n = |E|$, where the points of E are represented by vectors over $\mathbb{F}$. Assume that $\mathrm{pw}(M) = t$.

(I) For $i = 1, 2, \ldots, n$, suppose that we have got a sequence $X = (e_1, \ldots, e_{i-1}) \in E^{i-1}$ such that there exists a path-decomposition of M of width t which starts with the prefix X (note that initially $X = \emptyset$ and our assumption is trivial).

(II) For each $f \in E \setminus X$, we set $X_f = (e_1, \ldots, e_{i-1}, f)$. If $\lambda_M(X_f) \leq t$, we test whether there exists a path-decomposition of M of width t which starts with the prefix X_f.

(III) If the test of (II) succeeds for some (any) f—which has to happen for at least one value by the assumption—we let $X := X_f$ and continue with (I).

Clearly, this scheme results in the construction of a path-decomposition $(e_1, \ldots, e_n)$ of M of width t. The rest of the outline is devoted to explanation of crucial step (II).

It might seem that we could, in step (II), directly use the same approach from [8]. For convenience, we refer by X_f also to the underlying set of the sequence X_f. Indeed, let U be a copy of the uniform matroid $U_{k,2k}$, where $k = \lambda_M(X_f) \leq t$, such that all points of U are in the guts of the separation $(X_f, E \setminus X_f)$. Let $M' = (M \setminus X_f) \cup U$ as a union of point sets,

and $E' = E(M')$, $m = |E'|$. Then in any path-decomposition $(e'_1, \ldots, e'_m)$ of M' there is an index j such that $|U \cap \{e_1, \ldots, e_j\}| = k$, and hence the guts of the separation at position j contains all elements of a basis of U and so whole $\langle U \rangle$. However, this is not sufficient for our purpose—since we cannot branch a path-decomposition, we would additionally need U to form a prefix (or suffix) of the decomposition $(e'_1, \ldots, e'_m)$ which cannot be guaranteed in this way. Due to this setback, the case of path-width appears noticeably more difficult than that of branch-width (cf. also the discussion in [10]).

Our new solution is to, instead of just the uniform point set U of rank $k \leq t$, add to $M \setminus X_f$ a special set D of points of rank $t+1$ and path-width t. Let M'' denote the new matroid. The set D in M'', very informally saying, not only "fills in" the guts of the former separation $(X_f, E \setminus X_f)$ (which thus has to be displayed in any optimal path-decomposition of M''), but D also "fully occupies" one side (say, a prefix) of this guts in any optimal path-decomposition and so the elements of $E \setminus X_f$ must all occur after the guts. That is exactly what we need in step (II); the path-width of M'' is $\leq t$ if and only if the answer to (II) is Yes. Again, we use a suitable finite extension field of $\mathbb{F}$ to represent D. The details are given below, in Algorithm 3.1 and its proof.

▶ **Algorithm 3.1.** Let $\mathbb{F}$ be a fixed finite field and $t \in \mathbb{N}$ a fixed parameter. Let $\mathcal{P}$ be an oracle which, given any matroid N represented over $\mathbb{F}$, correctly decides whether $\mathrm{pw}(N) \leq t$. Let M be an input connected n-element matroid of rank r, given as an $r \times n$ matrix $\boldsymbol{A}$ over $\mathbb{F}$, and assume $\mathrm{pw}(M) = t$.

1. We pad $\boldsymbol{A}$ with 0's to make an $(r + t + 1) \times n$ matrix (adding "extra dimensions"). For simplicity, we will refer to the columns of the matrix as to the elements of M, with understanding that all computations will be carried out by means of linear algebra (i.e., in the matrix) in a natural way.
2. Let initially $X := \emptyset$. For $i = 1, 2, \ldots, n$, we repeat the following instructions:
   - **a.** We have got $X = (e_1, \ldots, e_{i-1}) \in E(M)^{i-1}$ where the elements of the sequence are distinct, and we use the symbol X to refer both to the sequence and the underlying set of elements of M.
   - **b.** We choose $f \in E(M) \setminus X$ such that $\lambda_M(X \cup \{f\}) \leq t$, and set $X_f := (e_1, \ldots, e_{i-1}, f)$.
   - **c.** We compute the guts $\Gamma := \langle X_f \rangle \cap \langle E(M) \setminus X_f \rangle$ and choose a subspace $\Sigma \supseteq \Gamma$ of rank exactly t and an element $d_0 \notin \Sigma$, such that $\langle \Sigma \cup \{d_0\} \rangle \cap \langle E(M) \rangle = \Gamma$ – informally, we use some of the "extra dimensions" from step (1) for placing Σ and d_0. Let P denote the set of all points of Σ (in the finite projective geometry over $\mathbb{F}$).
   - **d.** Let N_0 denote the matroid of rank $t+1$ induced by the points of $P \cup \{d_0\}$, and $\mathbb{F}_0 = \mathbb{F}$. For $j = 1, 2, \ldots, t$, let N_j be the matroid constructed as a free extension of N_{j-1} by an element d_j. As proved in Lemma 2.6, N_j is represented over the extension field $\mathbb{F}_j$ obtained from $\mathbb{F}_{j-1}$ by adjoining a root of degree $r(N_0) = t + 1$. At the end, let $D_0 := \{d_0, d_1, \ldots, d_t\}$, $D := P \cup D_0$ and $\mathbb{F}' = \mathbb{F}_t$.
   - **e.** For the matroid M' induced on the point set $(E(M) \setminus X_f) \cup D$ in the projective geometry over $\mathbb{F}'$, we ask the oracle $\mathcal{P}$ whether $\mathrm{pw}(M') \leq t$.
     - If the answer is No, then we repeat the steps from (2b) for another choice of f.
     - If the answer is Yes, then we update $X := X_f$ and continue the cycle in step (2) with the next value of i until $i = n$.
3. We output the path-decomposition $X = (e_1, \ldots, e_n)$ of M of width t.

Note that, in step 2e, some element e of M may be in the guts of $(X_f, E \setminus X_f)$ and then e is represented by the same point as some element of P in M'. It actually does not matter at all whether we consider these two elements as identical or a parallel pair.

▶ **Theorem 3.2.** *Let $\mathbb{F}$, t and $\mathcal{P}$ be as in Algorithm 3.1. For any connected n-element input matroid M represented by a matrix over $\mathbb{F}$, such that $\mathrm{pw}(M) = t$, Algorithm 3.1 correctly outputs a path-decomposition of M of width t. With fixed parameters $\mathbb{F}$ and t, the algorithm computes in FPT time $\mathcal{O}(n^4)$ and, in addition, makes $\mathcal{O}(n^2)$ calls to the oracle $\mathcal{P}$.*

Proof. We start with justifying correctness of the algorithm. Thanks to the condition $\lambda_M(X \cup \{f\}) \leq t$ in step (2b) of Algorithm 3.1, we know that the (eventual) output of the algorithm must be a path-decomposition of M of width t. Consequently, it is enough to prove that for every iteration of step (2) there is a choice of $f \in E(M) \setminus X$ which correctly succeeds in the test of step (2e). Assuming, for this moment, that it holds in step 2e;

(*) $\mathrm{pw}(M') \leq t$ if, and only if, there exists a path-decomposition of M of width t which starts with the prefix X_f,

the rest of the proof follows by a straightforward induction on i.

It is hence enough to prove the latter statement (*). In one direction ($\Leftarrow$), assume that there exists a path-decomposition $Y = (e_1, \ldots, e_n)$ of M of width t which starts with the prefix X_f. We can easily give a path-decomposition $Y' = (e'_1, e'_2, \ldots)$ of the matroid N_t induced by the point set D, where $e'_1 = d_0, e'_2 = d_1, \ldots, e'_{t+1} = d_t$ and this is followed by the elements of P in any order. The separation at position $j+1$ in Y', for $j < t$, has the guts $\langle\{d_0, \ldots, d_j\}\rangle$ of rank at most t. At positions $j+1$ for $j \geq t$, on the other hand, the guts is always Σ of rank t (or its subspace). We can now concatenate Y' with the suffix $(e_{i+1}, \ldots, e_n)$ of Y to form a path-decomposition of width t of the matroid M', since we have got $\langle D\rangle \cap \langle E(M)\rangle = \Gamma \subseteq \Sigma$ in step (2c). Therefore, $\mathrm{pw}(M') \leq t$.

In the opposite direction ($\Rightarrow$) of (*), assume that $\mathrm{pw}(M') \leq t$ and $Y' = (e'_1, \ldots, e'_p)$ is an optimal path-decomposition of M', where $p = |M'| = n - i + t + 1 + |P|$. Recall that P from step (2c) is the set of points of Σ over $\mathbb{F}$. It is a routine counting argument to show that for any $P_1 \cup P_2 = P$, $P_1 \cap P_2 = \emptyset$ and $|P_1| \leq |P_2| \leq |P_1| + 1$, the rank of both P_1, P_2 is t and so $\langle P_1\rangle \cap \langle P_2\rangle = \Sigma$. Hence there exists an index $1 \leq j \leq p$ such that the guts at position j in Y' contains Σ (and so it equals Σ and $\mathrm{pw}(M') = t$). Let $Y'_j = \{e'_1, \ldots, e'_j\}$ where $\lambda_{M'}(Y'_j) = t$ by the previous.

Recall the point set $D_0 = \{d_0, \ldots, d_t\}$ from step (2d). We first claim that, up to possible reversal of the sequence Y', it is $D_0 \subseteq Y'_j$. This easily follows from the conclusion of Lemma 2.3 since, for any $0 \leq a < b \leq d$, it is $\langle\Sigma \cup \{d_a\}\rangle = \langle\Sigma \cup \{d_b\}\rangle = \langle\Sigma \cup D_0\rangle$ of rank $t+1$, and so the condition of the lemma $\mu_{M'}(Y'_j, \{d_a\}) = \mu_{M'}(Y'_j, \{d_b\}) = \mu_{M'}(Y'_j, \{d_a, d_b\}) = t+1$ holds true.

Second, we claim that $(E(M) \setminus X_f) \cap Y'_j \subseteq cl_M(P)$ (geometrically, it is $cl_M(P) \subseteq \Sigma$). Suppose the contrary, that $Z := \big((E(M) \setminus X_f) \cap Y'_j\big) \setminus cl_M(P) \neq \emptyset$. Note that $Z \cap cl_M(P) = Z \cap \Sigma = \emptyset$. Under the assumption $Z \neq \emptyset$, we first consider the case that $\langle Z\rangle \cap \Sigma \neq \emptyset$. Informally, we are going to argue that the spans of Z and D_0 "freely overlap" in Σ and, for any $g \in Z \cup D_0$, the span of $(Z \cup D_0) \setminus \{g\}$ still contains $\langle D_0\rangle$. Then the path-decomposition Y', at some position before j, must contain $\langle D_0\rangle$ in the guts, but this is impossible since the rank of D_0 is $t+1$. The corresponding formal argument follows.

Let $M'' = M' \setminus D_0$. We choose $Z_0 \subseteq Z$ minimal by inclusion such that $\langle Z_0\rangle \cap \Sigma \neq \emptyset$, and so the rank of $\langle Z_0\rangle \cap \Sigma$ is one (in matroid terms this reads $r_{M''}(Z_0) + r_{M''}(P) = r_{M''}(Z_0 \cup P) + 1$). Since P contains all the points of Σ in the projective geometry over $\mathbb{F}$ (in matroid terms, P is a modular flat in M'' which is represented over $\mathbb{F}$), we have that $Z_0 \cap P \neq \emptyset$, and by minimality of Z_0 it is $Z_0 \cap P = \{p_0\}$. Consequently, $C_0 = Z_0 \cup \{p_0\}$ is a circuit in M'' and so also in M'. Now we look at the set $C_1 := D_0 \cup \{p_0\}$ in M' which is of rank $t+1$ and cardinality $t+2$, and hence is dependent. Since $d_1, \ldots, d_t$ have been chosen as free extensions

in step (2d), there cannot be any smaller circuits in C_1 and so C_1 itself is a circuit. We apply Lemma 2.2 to C_0 and C_1, obtaining a circuit $C_2 := C_0 \Delta C_1$ of M', where no element of C_2 belongs to Σ (our guts at position j). Since $C_2 \supseteq D_0$, the span of C_2 contains Σ and so $\mu_{M'}(Y'_j, C_2) = r_{M'}(C_2)$ and the conditions of Lemma 2.5 are fulfilled for C_2. However, the conclusion of the lemma contradicts our assumption $C_2 \subseteq Y'_j$.

Next, under the assumption $Z \neq \emptyset$, we consider the case that $\langle Z\rangle \cap \Sigma = \emptyset$. Since $\langle Z\rangle \cap \langle X_f\rangle \subseteq \Gamma \subseteq \Sigma$ by step (2c) of the algorithm, we have $\langle Z\rangle \cap \langle X_f\rangle = \emptyset$. Let $Z' = (E(M) \setminus X_f) \setminus Y'_j$. Since $\langle Z\rangle \cap \langle Z'\rangle \subseteq \Sigma$ from the path-decomposition Y' (at position j) of M', we similarly get $\langle Z\rangle \cap \langle Z'\rangle = \emptyset$. However, $\lambda_M(Z) \geq 1$ since M is connected, which means $\langle Z\rangle \cap \langle X_f \cup Z'\rangle \neq \emptyset$. Consequently, $\langle Z\rangle \cap \langle \Sigma \cup Z'\rangle \neq \emptyset$, but this in turn means, again from the path-decomposition Y' of M' with the guts Σ at position j, that $\langle Z\rangle \cap \Sigma \neq \emptyset$ – the case already being considered above.

To recapitulate (*), we have so far shown (by means of contradiction) that the path-decomposition Y' of M' has the guts Σ at position j, and $(E(M) \setminus X_f) \cap Y'_j \subseteq cl_M(P)$. Since $cl_M(P) \subseteq \langle P\rangle = \Sigma$ and Σ is the (geometric) guts of the separation $(X_f, E(M) \setminus X_f)$, we may now simply restrict Y' to the elements of $E(M) \setminus X_f$ and prefix this sequence with X_f, obtaining a valid path-decomposition of M of width t. The proof of (*) is finished.

The last point is to address runtime complexity of Algorithm 3.1. Note that the finite field $\mathbb{F}$ and the value of t are fixed parameters. In particular, arithmetic operations over $\mathbb{F}$ and $\mathbb{F}'$ (which depends only on $\mathbb{F}$ and t) take constant time each. Also note that $r \leq n$. We n times iterate at step (2), and each iteration costs the following. We are choosing at most n values of f in step (2b), and for each we compute the subspace Γ. Knowing $\langle X\rangle \cap \langle E(M)\setminus X\rangle$ already from the previous level, the computation of $\Gamma = \langle X_f\rangle \cap \langle E(M)\setminus X_f\rangle$ takes $\mathcal{O}(n^2)$ in step 2d by standard linear algebra (the rank of Γ is at most constant t). The rank of $\Sigma \supseteq \Gamma$ and cardinality of the set P are constants depending on $\mathbb{F}$ and t. Step (2d) takes $\mathcal{O}(1)$ time since it depends only on $\mathbb{F}$ and t and not on the input M. In fact, the point set D_0 needs to be computed only once during the whole algorithm and then linearly transformed to match actual Σ. This amounts to $\mathcal{O}(n^4)$ total time and $\mathcal{O}(n^2)$ calls to the oracle $\mathcal{P}$ in step (2e). ◀

4 Algorithmic Consequences

So far, in Section 3 we have restricted attention to connected matroids, but this is not any problem since we may easily concatenate path-decompositions of connected components of a general matroid. To make use of Algorithm 3.1, we also need to provide an implementation of the oracle $\mathcal{P}$ (which tests the value of path-width $\leq t$). This is significantly easier than computing an optimal path-decomposition thanks to Theorem 4.1. A class $\mathcal{N}$ of matroids is *minor-closed* if, for every matroid $M \in \mathcal{N}$, all minors of M also belong to $\mathcal{N}$. A matroid $M \notin \mathcal{N}$ is an *obstruction for membership in* $\mathcal{N}$ if all proper minors of M belong to $\mathcal{N}$.

▶ **Theorem 4.1** (Geelen–Gerards–Whittle [3],[1] and Hliněný [7]). *Let $\mathbb{F}$ be a fixed finite field and $k \in \mathbb{N}$ a fixed parameter. For any minor-closed class $\mathcal{N}$ of matroids, there are finitely many obstructions for membership in $\mathcal{N}$ which are representable over $\mathbb{F}$ and have branch-width at most k. Consequently, there is an FPT algorithm which, given an n-element matroid M represented by a matrix over $\mathbb{F}$, in time $\mathcal{O}(n^3)$ correctly decides whether $M \in \mathcal{N}$ or outputs that the branch-width of M is more than k.*

[1] We remark that Geelen, Gerards and Whittle have announced a "matroid minors" theorem which does not require a bound on branch-width to claim finite number of $\mathbb{F}$-representable obstructions for $\mathcal{N}$, but that is not fully published yet. For our purpose, the version of [3] is sufficient.

Direct implementation. The way we use Theorem 4.1 in an implementation of the oracle $\mathcal{P}$ combines Lemma 2.1 with Lemma 2.4; the matroids of path-width at most t have branch-width at most $t+1$ and form a minor-closed class $\mathcal{P}_t$ for which we can test membership in FPT time $O(n^3)$. Note, though, that this approach results in a *non-uniform FPT* algorithm (meaning that there is a sequence of algorithms for each value of the parameter t, rather than one universal algorithm), since we do not explicitly know the finite lists of obstructions for the classes $\mathcal{P}_t$, $t \in \mathbb{N}$. In combination with Theorem 3.2 we immediately get:

▶ **Corollary 4.2.** *Let $\mathbb{F}$ be a fixed finite field and $t \in \mathbb{N}$ a fixed parameter. There is a non-uniform FPT algorithm parameterized by t and $|\mathbb{F}|$ which, given an n-element matroid M represented by a matrix over $\mathbb{F}$, in time $\mathcal{O}(n^3)$ decides whether* $\mathrm{pw}(M) \leq t$.
Consequently, if $\mathrm{pw}(M) \leq t$, *there is a non-uniform FPT algorithm parameterized by t and $|\mathbb{F}|$, which in time $\mathcal{O}(n^5)$ outputs a path-decomposition of M of width t.* ◀

We remark that, in the setting of non-uniform algorithms, Algorithm 3.1 as used in Corollary 4.2 can be further simplified by the following observation. The point configuration D constructed in step (2d) is unique, up to a linear transformation, for given parameters $t, \mathbb{F}$, and hence it can be hard-coded into the (anyway non-uniform) algorithm with the smallest possible extension field $\mathbb{F}'$ which can represent D (this would quite likely be a much smaller field than the one computed by brute force in step (2d)).

Improving runtime. Runtime dependence on n of the algorithm of Corollary 4.2 can be improved to $\mathcal{O}(n^3)$ by using the same implementation tricks as in [8], based on earlier [6]. We briefly sketch this improvement which is again mainly of theoretical importance.

▶ **Theorem 4.3.** *Let $\mathbb{F}$ be a fixed finite field and $t \in \mathbb{N}$ a fixed parameter. There is a non-uniform FPT algorithm parameterized by t and $|\mathbb{F}|$ which, given an n-element matroid M represented by a matrix over $\mathbb{F}$, in time $\mathcal{O}(n^3)$ outputs a path-decomposition of M of width t or certifies that* $\mathrm{pw}(M) > t$.

Proof sketch. In the improved algorithm, we follow the general scheme of [8, Section 6] but in a simplified way. This is possible thanks to the fact that Algorithm 3.1, at each iteration, works with only one "active guts" of separation $(X, E \setminus X)$, unlike the algorithm of [8] which builds many branches of the desired branch-decomposition concurrently.

We modify the main steps of Algorithm 3.1 as follows:

1. For the input matroid M represented by the matrix $\boldsymbol{A}$, we use [6] to compute a branch-decomposition of M of width at most $3t+3$ —actually, a so-called $3t$-*boundaried parse tree* $\mathcal{T}$ for M —or to confirm that $\mathrm{bw}(M) > t+1$ and so $\mathrm{pw}(M) \geq \mathrm{bw}(M) - 1 > t$. This step takes $\mathcal{O}(n^3)$ time for fixed $t, \mathbb{F}$.
2. Each task performed in steps 2c and 2d can be done in time $\mathcal{O}(n)$ within the parse tree $\mathcal{T}$ (we refer to [8, Section 6] for corresponding details). It is important to compute within $\mathcal{T}$ (and not on whole $\boldsymbol{A}$), for which purpose we each time "enlarge" every node of $\mathcal{T}$ by the constant-rank subspace Σ. Subsequently, step 2e can test $\mathrm{pw}(M') \leq t$ by checking (non-)presence of the finitely many obstructions for 'path-width $\leq t$'. This test can also be done in time $\mathcal{O}(n)$ by [7] since minor obstructions are MSO-definable.

Altogether, runtime is $\mathcal{O}(n^3 + n^2 \cdot n) = \mathcal{O}(n^3)$ for fixed $t, \mathbb{F}$. ◀

Abstract matroids. Besides its simplicity, our Algorithm 3.1 has another important theoretical advantage over the constructive algorithm of [10]. While [10] directly compute with points and subspaces in a finite projective geometry, and it does not seem possible to extend

their approach to infinite projective geometries or abstract matroids, we can easily adapt our algorithm to work even with *abstract matroids* given by a rank oracle (although our algorithm also directly worked with the points of a subspace Σ, that was only for convenience and clarity, and could be rather easily replaced by an abstract handling).

In this respect we mention the algorithm of Nagamochi [13] which computes an optimal path-decomposition for an arbitrary submodular function (and hence including the case of a matroid given by a rank oracle). Moreover, the algorithm [13] is self-contained in the sense that it does not call an external decision subroutine (as we are going to do here). Though, its runtime is of order $\mathcal{O}(n^{f(t)})$ where t is the path-width (complexity class XP) while we aim for an FPT algorithm, that is with runtime of order $\mathcal{O}(f(t)\cdot n^c)$.

We say that an abstract matroid M is given by a *rank oracle* if the input consists of the ground set $E = E(M)$ and an oracle function $\mathcal{R} : 2^E \to \mathbb{N}$ such that $\mathcal{R}(X) = r_M(X)$ for all $X \subseteq E$. Algorithms then handle M by asking $\mathcal{R}$ so called *rank queries*.

▶ **Theorem 4.4.** *Let $t \in \mathbb{N}$ and $\mathcal{P}$ be an oracle function which, for any matroid N given by a rank oracle, correctly decides whether* $\mathrm{pw}(N) \le t$. *There is an algorithm that, for an input n-element matroid M given by a rank oracle $\mathcal{R}$, outputs a path-decomposition of M of width t or correctly answers that* $\mathrm{pw}(M) > t$. *The algorithm makes $\mathcal{O}(n^2)$ calls to the oracle function $\mathcal{P}$ and, neglecting the fixed parameter t, asks $\mathcal{O}(n^2)$ rank queries.*

Before moving onto the proof, we need one more technical concept. We are going to modify the matroid M (which we do not completely know—we cannot read all the ranks of sets in M in polynomial time!). Instead, we will modify the rank oracle $\mathcal{R}$ by prescribing its (efficient) answers to rank queries involving elements which we add to M. In this respect we define the following three elementary operations:

(M1) *Adding a coloop a* to M defines, for every $X \subseteq E(M)$, that $\mathcal{R}(X \cup \{a\}) := \mathcal{R}(X) + 1$.
(M2) *Placing b freely into the closure of $Z \subseteq E(M)$* defines, for every $X \subseteq E(M)$,
- $\mathcal{R}(X \cup \{b\}) := \mathcal{R}(X)$ if $r_M(Z \cup X) = r_M(X)$, and
- $\mathcal{R}(X \cup \{b\}) := \mathcal{R}(X) + 1$ otherwise.

(M3) *Placing c freely into the guts of* (the separation of) $Z \subseteq E(M)$ means, for $X \subseteq E(M)$,
- $\mathcal{R}(X \cup \{c\}) := \mathcal{R}(X)$ if $\mu_M(Z, X) = r_M(X)$, and
- $\mathcal{R}(X \cup \{c\}) := \mathcal{R}(X) + 1$ otherwise.

An informal geometric explanation of these operations follows. (M1) simply "adds another dimension" with a. (M2) puts the new point b in general position (i.e., without unforced linear dependencies) into the span $\langle Z\rangle$. (M3) similarly puts the new point c in general position into the guts $\langle Z\rangle \cap \langle E(M) \setminus Z\rangle$. It is a routine exercise to prove that the rank oracle defined by each one of (M1), (M2), (M3) is the rank function of a matroid.

Proof sketch of Theorem 4.4. Again, we may restrict our attention to connected input matroids M. We modify some steps of Algorithm 3.1 as follows:

- Step (1) is not needed.
- In step (2c), let $k = \lambda_M(X_f)$. First, we $(t-k)$-times (if $k < t$) repeat the operation (M1) of adding a coloop. Let P_0 denote the set of coloops added to M this way. We then $(t+k)$-times repeat the operation (M3) of placing a new element freely into the guts of $\big(X_f \cup P_0,\, (E(M) \setminus X_f) \cup P_0\big)$—to be formally precise, we consider for this operation the elements of P_0 duplicated. Let $P \supseteq P_0$ denote the set of all the $2t$ added elements, which is of rank t (one may observe that P actually induces a uniform matroid $U_{t,2t}$).

- In step (2d), we add a new coloop d_0 by (M1). Then, for $j = 1, \ldots, t$, we iteratively do the operation (M2) of freely placing a new element d_j into the closure of $P \cup \{d_0\}$. Again, let $D_0 := \{d_0, d_1, \ldots, d_t\}$ and $D := P \cup D_0$.
- In step (2e), we let M' to be the matroid defined on the ground set $(E(M) \setminus X_f) \cup D$ by the rank oracle $\mathcal{R}'$ constructed from $\mathcal{R}$ by the above modifications.

In the proof of the modified algorithm, we can essentially repeat nearly all the arguments of the proof of Theorem 3.2, translated into the abstract setting of the rank functions of M and M'. There are, though, two points in the proof of the forward direction of (*) which need different arguments:

- First, assuming a path-decomposition $Y' = (e'_1, \ldots, e'_p)$ of M' of width t, we define index j as the minimum $1 \leq j \leq t$ such that $|P \cap Y'_j| = t$. Since the elements of P have been each freely placed into a rank-t flat, the t-element set $P \cap Y'_j$ is independent, and so is the complement $P \setminus Y'_j$. Consequently, all elements of P belong to the guts of the separation at the position j of Y', analogously to the former proof.
- Second, assuming a nonempty set $Z := \big((E(M) \setminus X_f) \cap Y'_j\big) \setminus cl_M(P) \neq \emptyset$, we have to find a circuit $C_2 \subseteq Y'_j$ which then provides a contradiction to Lemma 2.5. We can again argue by connectivity of M that $r_M(Z) + r_M(P) > r_M(Z \cup P)$. Since $cl_{M'}(D_0) \supseteq P$, the set $Z \cup D_0$ is dependent in M', and one can derive that, since D_0 is independent and its elements are freely placed in the closure (flat) of D_0, a circuit in $Z \cup D_0$ must contain whole D_0 and a suitable subset $Z_0 \subseteq Z$, giving desired $C_2 := D_0 \cup Z_0$.

We skip further technical details in the restricted short paper (see the Appendix). ◀

5 Conclusions

We have shown a relatively simple oracle algorithm which can construct an optimal path-decomposition of a given matroid if it is provided with a subroutine testing the value of matroid path-width. This provides a significantly simpler alternative to the complicated SODA'16 algorithm of Jeong, Kim and Oum [10] which uses a direct construction based on ideas originally developed for graphs by Bodlaender and Kloks [1]. Though, there is price we have to pay for simplicity; our approach provides a non-uniform FPT algorithm, caused by the fact that we have no explicit bound on the size of the minor-minimal obstructions for path-width $\leq t$ (unlike the case of branch-width in which an explicit bound [2] readily provides a uniform FPT algorithm [8]).

Comparing our result to Jeong, Kim and Oum [10], there are two points on which we should comment. Our approach currently cannot handle path-decompositions of generally partitioned matroids as [10] (as well as older [8]) do. The problem actually lies in the definition of path-width which does not restrict the rank of each one part. Our algorithm, when generalized to partitioned matroids as in [8], would compute an optimal linear branch-decomposition which is, for partitioned matroids in general, a different notion than a path-decomposition. Though, in particular, our algorithm *can* straightforwardly compute, using the same reduction as in [8], an optimal linear rank-decomposition of a graph as [10] does.

On the other hand, unlike [10], our algorithm readily generalizes to abstract matroids given by rank oracles, as proved in Theorem 4.4. Although we are currently not aware of an FPT algorithm which could test path-width $\leq t$ for matroids given by rank oracles, such algorithms could probably emerge in the future (cf. [13]) for other matroid classes, and then Theorem 4.4 will be readily applicable also to these new classes. Along the same line, it is likely that in the future an explicit bound on the obstructions for path-width $\leq t$ will be found and then our FPT algorithm would immediately become uniform.

References

1 H. L. Bodlaender and T. Kloks. Efficient and constructive algorithms for the pathwidth and treewidth of graphs. *J. Algorithms*, 21(2):358–402, 1996.

2 J.F. Geelen, A.M.H. Gerards, N. Robertson, and G. Whittle. On the excluded minors for the matroids of branch-width *k*. *J. Comb. Theory, Ser. B*, 88(2):261–265, 2003.

3 J.F. Geelen, A.M.H. Gerards, and G. Whittle. Branch-width and well-quasi-ordering in matroids and graphs. *J. Comb. Theory, Ser. B*, 84(2):270–290, 2002.

4 J.F. Geelen, B. Gerards, and G. Whittle. On rota's conjecture and excluded minors containing large projective geometries. *J. Comb. Theory, Ser. B*, 96(3):405–425, 2006.

5 P. Hliněný. On matroid properties definable in the MSO logic. In *Mathematical Foundations of Computer Science, MFCS 2003, Proceedings*, volume 2747 of *Lecture Notes in Computer Science*, pages 470–479. Springer, 2003.

6 P. Hliněný. A parametrized algorithm for matroid branch-width. *SIAM J. Comput.*, 35(2):259–277, 2005.

7 P. Hliněný. Branch-width, parse trees, and monadic second-order logic for matroids. *J. Combin. Theory Ser. B*, 96(3):325–351, 2006.

8 P. Hliněný and S. Oum. Finding branch-decomposition and rank-decomposition. *SIAM J. Comput.*, 38:1012–1032, 2008.

9 K. Jain, I.I. Mandoiu, and V.V. Vazirani. The 'art of trellis decoding' is computationally hardi - for large fields. *IEEE Trans. Information Theory*, 44(3):1211–1214, 1998.

10 J. Jeong, E.J. Kim, and S. Oum. Constructive algorithm for path-width of matroids. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 1695–1704. SIAM, 2016.

11 N. Kashyap. Matroid pathwidth and code trellis complexity. *SIAM J. Discrete Math.*, 22(1):256–272, 2008.

12 A. Koutsonas, D.M. Thilikos, and K. Yamazaki. Outerplanar obstructions for matroid pathwidth. *Discrete Mathematics*, 315-316:95–101, 2014.

13 H. Nagamochi. Linear layouts in submodular systems. In *Algorithms and Computation - 23rd International Symposium, ISAAC 2012, Proceedings*, volume 7676 of *Lecture Notes in Computer Science*, pages 475–484. Springer, 2012.

14 James G. Oxley. *Matroid Theory (Oxford Graduate Texts in Mathematics)*. Oxford University Press, Inc., New York, NY, USA, 2006.

15 N. Robertson and P.D. Seymour. Graph minors. i. excluding a forest. *J. Comb. Theory, Ser. B*, 35(1):39–61, 1983.

16 A. Vardy. Trellis structure of codes. In *Handbook of coding theory. Vol. 1. Part 1: Algebraic coding. Vol. 2. Part 2: Connections, Part 3: Applications*, pages 1989–2117. Amsterdam: Elsevier, 1998.

A Appendix to Section 2

▶ **Lemma A.1** (Lemma 2.4). *Let M be a matroid and N a minor of M. Then* $\mathrm{pw}(N) \leq \mathrm{pw}(M)$.

Proof. Consider $X \subseteq E(M)$ and $e \in E(M)$. It is folklore that $\lambda_{M\setminus e}(X \setminus \{e\}) \leq \lambda_M(X)$ and $\lambda_{M/e}(X \setminus \{e\}) \leq \lambda_M(X)$. Hence, by induction on $|E(M)| - |E(N)|$, the restriction of any path-decomposition of M is a path-decomposition of N of at most the same width. ◀

▶ **Lemma A.2** (Lemma 2.5). *Let M be an n-element matroid and $(e_1, \ldots, e_n)$ be its path-decomposition of width $t = w_M(e_1, \ldots, e_n)$. For an index i let $X = \{e_1, \ldots, e_i\}$ and $Y = \{e_{i+1}, \ldots, e_n\} = E(M) \setminus X$ such that $\lambda_M(X) = t$. Assume that there exists a circuit $C \subseteq E(M)$ such that no element of C is in the guts of (X, Y) and $\mu_M(X, C) = r_M(C)$. Then $X \cap C \neq \emptyset \neq Y \cap C$.*

(The condition $\mu_M(X, C) = r_M(C)$ can be seen as that "C spans the guts of (X, Y).")

Proof. We proceed by means of contradiction, aiming to show that $w_M(e_1, \ldots, e_n) > t$. Up to symmetry, let $Y \cap C = \emptyset$, meaning that $C \subseteq X$. Let $j \leq i$ by the largest index such that $e_j \in C$, and $X' = \{e_1, \ldots, e_{j-1}\}$, $Y' = \{e_j, \ldots, e_n\}$. From the assumptions $\lambda_M(X) = r_M(X) + r_M(Y) - r(M) = t$ and $\mu_M(X, C) = r_M(X) + r_M(Y \cup C) - r(M) = r_M(C)$, we derive

$$r_M(C) + r_M(Y) - r_M(Y \cup C) = t. \tag{1}$$

It is $r_M(C \setminus \{e_j\}) = r_M(C)$ since C is a circuit, and $r_M(Y \cup \{e_j\}) = r_M(Y) + 1$ since e_j is not in the guts of the separation (X, Y). Hence we can rewrite (1) as

$$r_M(C \setminus \{e_j\}) + r_M(Y \cup \{e_j\}) - r_M(Y \cup C) = t + 1. \tag{2}$$

Note that $Y \cup \{e_j\} \subseteq Y'$ and $C \setminus \{e_j\} \subseteq X'$. We conclude the proof by showing

$$\begin{aligned}
\lambda_M(X') &= r_M(X') + r_M(Y') - r(M) \\
&= r_M(X') + \big(r_M(Y') - r_M(X' \cup Y')\big) \\
&\geq r_M(X') + \big(r_M(Y \cup \{e_j\}) - r_M(X' \cup Y \cup \{e_j\})\big) \\
&= r_M(Y \cup \{e_j\}) + \big(r_M(X') - r_M(X' \cup Y \cup \{e_j\})\big) \\
&\geq r_M(Y \cup \{e_j\}) + \big(r_M(C \setminus \{e_j\}) - r_M((C \setminus \{e_j\}) \cup Y \cup \{e_j\})\big) \\
&= r_M(C \setminus \{e_j\}) + r_M(Y \cup \{e_j\}) - r_M(C \cup Y) = t + 1,
\end{aligned}$$

using submodularity and (2). This however contradicts $w_M(e_1, \ldots, e_n) = t$. ◀

B Appendix to the proof of Theorem 4.4

In the last step of the proof of Theorem 4.4, right before deriving a contradiction to Lemma 2.5, we actually use a statement of forthcoming Lemma B.2 (its notation corresponds to that of Theorem 4.4 except that $X = X_f$). The appendix is devoted to a formal proof of this statement. First, supplementary Lemma B.1 is given.

▶ **Lemma B.1.** *Let M be a matroid and $Y \subseteq E = E(M)$, $Y' = E \setminus Y$. Assume that $Q \subseteq E$ is such that all elements of Q are in the guts of the separation (Y, Y') and $r_M(Q) = \lambda_M(Y)$. If $C \subseteq E$ is a circuit and $e \in Y' \cap C$, then there exists a circuit C' such that $e \in C' \subseteq (C \setminus Y) \cup Q$ and $C' \subseteq cl_M(Y')$.*

Proof. If $C \subseteq cl_M(Y')$, we are done. Otherwise, it is $C \cap Y' \neq C$ is independent and $r_M(C \cap Y) + r_M(C \cap Y') - r_M(C) > 0$ since a circuit is connected, and so $r_M(C \cap Y') + r_M(Y) - r_M(C \cup Y) > 0$ by submodularity. Let $Q_1 \subseteq Q$ be a basis of Q. Note that $Q_1 \subseteq cl_M(Y')$. Since Q_1 "spans the guts" of (Y, Y') (formally, $r_M(Q_1) = \lambda_M(Y')$), the set $C_1 = (C \cap Y') \cup Q_1$ is dependent and contains a circuit C'. Since Q_1 is independent, we may choose C' such that $e \in C'$. Finally, $C' \subseteq (C \cap Y') \cup Q = (C \setminus Y) \cup Q$ and $C' \subseteq cl_M(Y')$ are true by the definition of C'. ◀

▶ **Lemma B.2.** *Let M be a connected matroid, $X \subseteq E = E(M)$ and $\emptyset \neq Z \subseteq E \setminus X$. Assume that $P \subseteq E \setminus X$ is such that $r_M(P) = t \geq \lambda_M(X)$, $\mu_M(X, P) = t$, and $Z \cap cl_M(P) = \emptyset$. Furthermore, assume that M_0 is a matroid on the ground set $E \cup D_0$ where $D_0 = \{d_0, d_1, \ldots, d_t\}$, the restriction of M_0 to E is M, and d_0 is a coloop w.r.t. M and each d_i is freely placed in the closure of $P \cup \{d_0\}$ w.r.t. $E \cup \{d_0, \ldots, d_{i-1}\}$ for $i = 1, \ldots, t$. Let M' be M_0 restricted to $E(M_0) \setminus X$. If $\lambda_{M'}(D_0 \cup Z) \leq t$, then there exists $Z_0 \subseteq Z$ such that $D_0 \cup Z_0$ is a circuit of M'.*

Proof. First, note that D_0 is independent both in M_0 and M' by the construction, of rank $t+1$ and $cl_{M'}(D_0) \supseteq P$, but $D_0 \cap cl_M(P) = \emptyset$. From the assumption $\lambda_{M'}(D_0 \cup Z) \leq t$ we also see that P "spans" the guts of $(D_0 \cup Z, E \setminus Z)$ in M_0, i.e., $r_{M_0}(P) = t = \lambda_{M_0}(D_0 \cup Z) = \lambda_{M'}(D_0 \cup Z)$ and all elements of P are in this guts (both in M_0 and in M').

We are going to apply twice the previous lemma. It is a folklore fact that in a connected matroid (which is a property analogous to 2-connected graphs), every two elements belong to a common circuit. Let $e_1 \in X$ and $e_2 \in Z$ be arbitrary and $C \ni e_1, e_2$ be a circuit of M. We apply Lemma B.1 to M, C and $Y := X$, $Q := P$, $e := e_2$. The resulting circuit C' satisfies: $e_2 \in C' \subseteq (C \setminus X) \cup P$ and $C' \subseteq cl_M(E \setminus X)$.

Then, in the matroid M', we let $Y := cl_{M'}((E \setminus X) \setminus Z)$ where, clearly, $P \subseteq Y$ and $Z \cap Y = \emptyset$ since $Z \cap cl_{M'}(P) = \emptyset$. In the matroid $M'' = M'/d_0$ obtained by contracting d_0, it holds $D_1 = \{d_1, \ldots, d_t\} \subseteq cl_{M''}(P)$ and $cl_{M''}(D_1) = cl_{M''}(P)$, since $cl_{M'}(D_0) \supseteq P$ and $r_{M''}(D_1) = t = r_{M''}(P)$. Note that C' is a circuit of M'', too, since M'' restricted to $E \setminus X$ equals $M \setminus X$. We now apply Lemma B.1 to M'', Y and $Q := D_1$, $C := C'$, where $e_2 \in Z \cap C' \subseteq Y' \cap C'$. The circuit C'' that we obtain, satisfies $C'' \subseteq (C' \setminus Y) \cup D_1 = Z \cup D_1$.

Let $Z_0 := Z \cap C''$. Back in M' (uncontracting d_0), $C'' \cup \{d_0\}$ is a circuit of M', and $C'' \supseteq D_1$ since the elements of D_1 have been freely placed—they do not have unforced dependencies in M'. Hence this circuit is $D_0 \cup Z_0 = C'' \cup \{d_0\}$. ◀